arXiv:1610.05857v1 [stat.ME] 19 Oct 2016

GRAPHICAL MODELS FOR ZERO-INFLATED SINGLE CELL GENE EXPRESSION

By Andrew McDavid*,‡ Raphael Gottardo*,‡ Noah Simon† and Mathias Drton‡

Department of Statistics, Department of Biostatistics†, University of Washington; Vaccine and Infectuous Disease Division‡, Fred Hutchinson Cancer Research Center, Seattle, Washington*

Bulk gene expression experiments relied on aggregations of thousands of cells to measure the average expression in an organism. Advances in microfluidic and droplet sequencing now permit expression profiling in single cells. This study of cell-to-cell variation reveals that individual cells lack detectable expression of transcripts that appear abundant on a population level, giving rise to zero-inflated expression patterns. To infer gene co-regulatory networks from such data, we propose a multivariate Hurdle model using a finite mixture of singular Gaussian distributions. This permits inference of statistical independences in zero-inflated, semi-continuous data to learn undirected, Markov graphical models. The node-wise conditional log-likelihood of the multivariate Hurdle model is convex and tractable, and allows neighborhood selection. We apply penalized maximum pseudo-likelihood using a group lasso penalty to infer conditional independences. The method is demonstrated in a data set of selected expression in T follicular helper cells, and a high-dimensional profile of mouse dendritic cells. It reveals network structure not present using other methods; or in bulk data sets.

1. Background. Graphical models have synthesized high-dimensional biological experiments into understandable, canonical forms [Dobra et al., 2004, Markowetz and Spang, 2007]. Although inferring causal relationships between genes is perhaps the ultimate goal of such analysis, causal models are difficult to estimate with observational data, while experimental manipulation of specific genes has remained costly, and largely inimitable to high-throughput biology. On the other hand, undirected models parametrize the conditional independences present between variables with a graph consisting of a set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. These Markov random fields are estimable on *iid* observational data, and provide descriptions of the statistical predictors of each gene. Each gene is optimally predicted using only its neighbors.

Improved descriptions of conditional independence in gene expression experiments would help answer a variety of scientific questions. They might

provide new insights on–or at least falsify–models of gene regulation, since statistical dependence is expected, given causal dependence. In immunology, polyfunctional immune cells, which simultaneously express multiple cytokines, have been identified as useful predictors of vaccine response [Precopio et al., 2007]. Simultaneous expression or *co-expression* of cellular surface markers has been used to define cellular phenotypes [Lin et al., 2015], so expanding the "dictionary" of co-expression may allow phenotypic refinement.

For example, T follicular helper (Tfh) cells are a class of $CD4^+$ lymphocytes. B-cells that secrete antibodies require Tfh cell co-stimulation to become active [Ma et al., 2012]. Tfh cells are defined, and identified both through their location in the B-cell germinal centers, as well as their production of high levels of the proteins CXCR5, PD1 and BcL-6. In an experiment detailed in section 6, Tfh cells were identified from $CD4^+CXCR5^+PD1^+$ cells from lymph node biopsy.

In figure 1 the expression of four Tfh marker genes ($P < 10^{-20}$ compared to non-Tfh lymph node T-cells; not shown) is plotted. Although the expression of these genes could help discriminant Tfh from non-Tfh cells, the strength of their linear correlations within Tfh cells (upper panels) varies. To identify coexpressing subsets of cells or to clarify the conditional relationship between genes, estimating the statistical dependences within Tfh cells is necessary. Also noteworthy is the zero-inflation of this single cell data. This causes the pairs of linear regression models (shown in the lower panels) to fit poorly.

Several different approaches have been described to estimate the structure of undirected graphical models [Drton and Maathuis, 2017]. Fully parametric joint models have assumed a Gaussian distribution [Yuan and Lin, 2007], or that the marginal distributions are monotone transformations thereof [Liu et al., 2009]. Pseudo-likelihood based models, which posit only the conditional distribution of each gene without necessarily guaranteeing joint compatibility [Arnold and Press, 1989, Meinshausen and Bühlmann, 2006, Chen et al., 2015] are more flexible, and have also allowed high-dimensional generalized additive models [Voorman et al., 2014], where the conditional expectation is a smooth, additive function. Many of these methods have seen profitable application to gene expression experiments on bulk aggregates of cells assayed through microarrays or RNA sequencing.

Microfluidic and molecular barcoding advances have enabled the measurement of the minute quantities of mRNA present in single cells. A characteristic of single cell expression is zero-inflation of otherwise continuous measurements, in which measurements are either strongly positive, or undetectable. These experiments have the potential to provide unique resolution

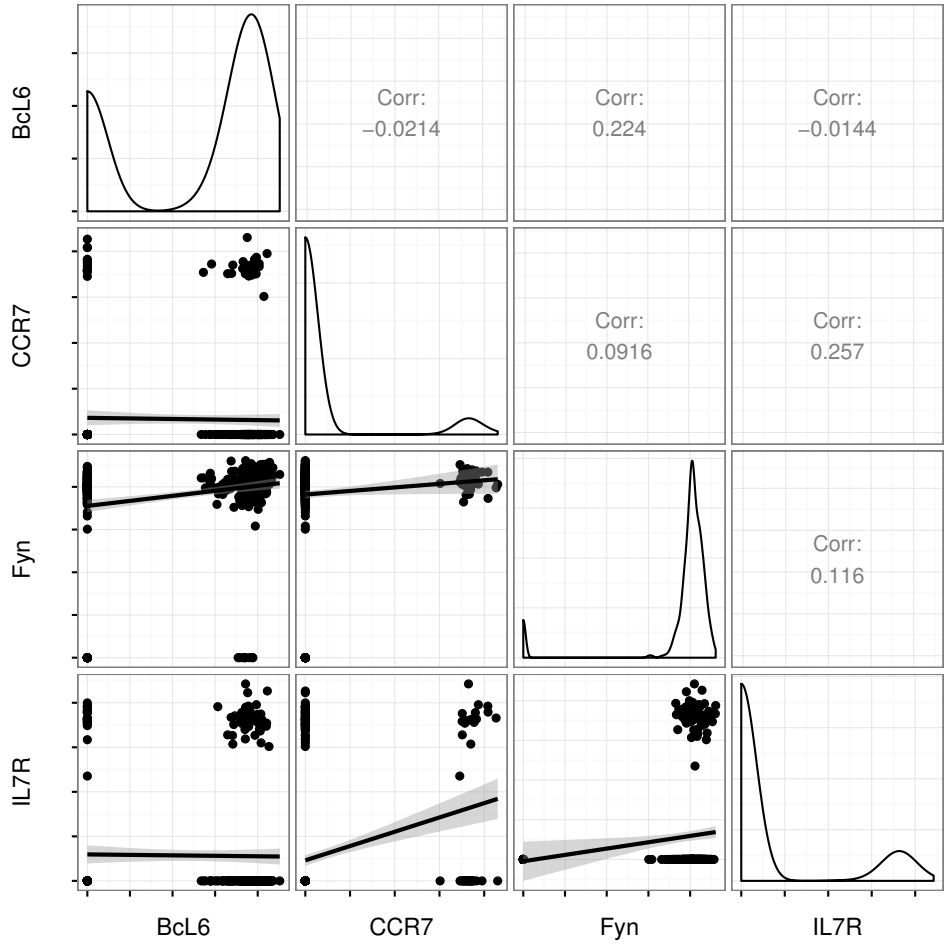

Fig 1: Scatter plots of inverse cycle threshold ($40 - \text{Ct}$) measurements from a quantitative PCR (qPCR)-based single cell gene expression experiment (lower panels). The cycle threshold (Ct) is the PCR cycle at which a predefined fluorescence threshold is crossed, so a larger inverse cycle threshold corresponds to greater log-expression [McDavid et al., 2013]. Measurements that failed to cross the threshold after 40 cycles are coded as 0.
Marginal expression in Tfh ($\text{CXCR5}^+\text{PD1}^+$) cells of Tfh marker genes is illustrated in the kernel-density estimates along the diagonal. The linear relationship between pairs of genes is shown in the lower panels.

of gene co-expression, but the distributions are inadequately modeled with a Gaussian distribution, and in any case the properties of the zero-inflation may be of intrinsic interest [Kim and Marioni, 2013]. To accommodate these features, we propose a joint probability density function $f(\mathbf{y})$ of the form

$$\log f(\mathbf{y}) = \mathbf{v_y}^T \mathbf{G} \mathbf{v_y} + \mathbf{v_y}^T \mathbf{H} \mathbf{y} - \frac{1}{2} \mathbf{y}^T \mathbf{K} \mathbf{y} - C(\mathbf{G}, \mathbf{H}, \mathbf{K}), \tag{1}$$

in which both binary and continuous versions of gene expression are sufficient statistics, and interactions thereof are parametrized. Here $\mathbf{y}$ represents an m-vector of gene expression, $[\mathbf{v_y}]_i = I_{y_i \neq 0}$ is the element-wise non-zero indicator function, and $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{K}$ are matrices of interaction parameters. This model can be shown equivalent to a finite mixture model of singular Gaussian distributions. The neighborhood of each gene can be estimated using an **anisometric** group-ℓ_1 penalized conditional likelihood of the form

$$\text{argmax}_\theta \log f_{[b|A]}(y; \theta) - \lambda \sum_{a \in A} \sqrt{\theta_a^T \mathbf{H}_{aa} \theta_a},$$

where $\log f_{[b|A]}(y; \theta)$ is the conditional likelihood of $y_b | \mathbf{y}_A$, θ is the concatenation of rows and columns of interaction matrices, and $\lambda \geq 0$ is a tuning parameter. Typically the group-ℓ_1 penalty [Yuan and Lin, 2006] takes $\mathbf{H}_{aa} = I$. We propose to use the observed Fisher information in block a under a null model $\theta_a = 0$ for all $a \in A$.

Section 2 of this paper discusses the parameter targeted in single cell gene expression experiments, and why it is not accessible from traditional gene expression experiments. Section 3 describes the parametric Hurdle model for single cell gene expression and estimation of graphical models using neighborhood selection via penalized regression. Section 4 provides a simulation study. Section 5 illustrates the method in a data set in which selected gene profiles were available for both single- and several-cell aggregates, while section 6 applies the method to a high-dimensional data set. The proposed method yields substantial improvements in simulation, and uncovers distinct networks compared to existing approaches.

2. From single cells to cellular co-expression. A typical cell contains 1-50 picograms of total RNA, of which perhaps 5% is assayable messenger RNA encoding for proteins (the remainder is structural tRNA and rRNA) [Livesey, 2003]. Protocols for bulk gene expression experiments, such as for Illumina TrueSeq, may call for 100 nanograms of total mRNA, hence require the equivalent of 80,000 cells' worth of mRNA. On the one hand, this biological "summation" over thousands of cells is expected to yield sharper

inference on the mean expression level of each gene. However, this comes at the cost of distorting the conditional dependences present between genes.

Consider $\mathbf{Y}_i$, an *iid* sequence of random vectors in $\mathbb{R}^p$ representing the copy number of p transcripts present in single cells $i = 1, \ldots, n$. Now suppose the n cells are aggregated and the total expression is measured using some linear quantification that reports values proportional to the input counts of mRNA. Then the sum of expression observed in *bulk* experiments is

$$\mathbf{Z} = \sum_{i}^{n} \mathbf{Y}_i.$$

Although most bulk experiments are designed to test for differences in mean expression due to experimental treatments and lack extensive replication within a condition, *stochastic profiling* [Janes et al., 2010] experiments have provided *iid* replicates of $\mathbf{Z}$ suitable for estimating higher order moments. But when the distribution of $\mathbf{Y}_i$ obeys some conditional independence relationships, in general the distribution of $\mathbf{Z}$ does not obey these same relationships. For example, take $p = 3$ and suppose that $\mathbf{Y}_i$ are *iid* samples from a tri-variate distribution $[Y_1, Y_2, Y_3]$ on $\{0, 1\}^3$. Suppose the probability mass function (PMF) factors as $p(y_1, y_2, y_3) = p(y_1)p(y_2|y_1)p(y_3|y_2)$, that is Y_1 and Y_3 are conditionally independent given Y_2, or in symbols $Y_1 \perp Y_3|Y_2$. Then $p(y_3, y_2|y_1) = p(y_2|y_1)p(y_3|y_2)$, which is equivalent to saying that each 2×2 probability table $p(y_3, y_2|y_1 = j)$, $j = 0, 1$ has non-negative rank one. Yet even summing over $n = 2$ cells, the PMF of $\mathbf{Z} = \mathbf{Y_1} + \mathbf{Y_2}$ will not generally factor as such, as one may exhibit a 3×3 probability table for $p(z_3, z_2|z_1)$ that is not a non-negative rank one matrix.

The infamous case in which graphical structure commutes under convolution is when the $\mathbf{Y}_i$ are multivariate Normal. But as argued in the next section, single cell gene expression is generally bimodal and zero-inflated, so not plausibly described by a multivariate Normal distribution. Even though for large enough n the distribution of the bulk experiment $\mathbf{Z}$ might approach multivariate (log-)normality, the networks estimated in bulk data will not reflect the conditional independences that hold in single cell data.

In the limit of n large, $\mathrm{Cov}(\mathbf{Z})$ converges to the population covariance of $\mathrm{Cov}(\mathbf{Y})$. In some models, and some graphs, such as tree-structured Ising models, the independence structure of $\mathbf{Y}$ can be at least partially identified from $\mathrm{Cov}(\mathbf{Y})^{-1}$, but generally there is no easy connection between $\mathrm{Cov}(\mathbf{Y})$ and the conditional independence structure of $\mathbf{Y}$ [Loh and Wainwright, 2013].

2.1. *Single cell expression.* A distinctive feature of single cell gene expression data–across methods and platforms–is the bimodality of expression values [McDavid et al., 2013, Finak et al., 2015, Marinov et al., 2014, Shalek et al., 2014]. Genes can be on (and a positive expression measure is recorded) or off–or below a limit of detection–and the recorded expression is zero or negligible. The cause of the distributional bimodality remains unresolved. It has been argued that it represents censoring of expression below a substantial limit of detection, yet comparison of *in silico* signal summation from many single cells, to the signal measured in biological sums of cells suggest that the limit of detection is essentially zero.

Moreover, the empirical distribution of the log-transformed counts appears rather different than would be expected from censoring: the distribution of the log-transformed, positive values is generally symmetric. Yet the presence of bimodality in technically replicated experiments ("Pool/split" experiments) implicates technical factors as a factor [Marinov et al., 2014].

2.2. *Markov graphs.* A Markov random field encodes the conditional independence between components of a random vector $\mathbf{Y}$ through a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E} \subseteq \mathcal{V} \times \mathcal{V})$. The set $\mathcal{V}$ contains vertices indexing $\mathbf{Y}$, while $\mathcal{E}$ is the edge set. This graph is undirected in the sense that if $(a, b) \in \mathcal{E}$ then so is (b, a). Associated with $\mathcal{G}$ is a probability distribution $P(\mathbf{Y})$. The graph specifies minimal forms of conditional independence on $P(\mathbf{Y})$.

Of primary use in this paper will be when $P(\mathbf{Y})$ is compatible with $\mathcal{G}$ under the **local Markov property**. A distribution $P(\mathbf{Y})$ satisfies this property with respect to $\mathcal{G}$ when for all $a \in \mathcal{V}$, $Y_a \perp Y_{\mathcal{V}\setminus \mathrm{cl}(a)} | Y_{\mathrm{ne}(a)}$, where $\mathrm{ne}(a)$ denotes the neighborhood (adjacent vertices) of a and $\mathrm{cl}(a) = \mathrm{ne}(a) \cup a$ is the closure. When $P(\mathbf{Y})$ has a positive and continuous density f with respect to some product measure, as is the case to be considered, then the local Markov property is equivalent to the pairwise and global Markov properties. The Hammersley-Clifford theorem then states that any density–and only these densities–that are compatible with $\mathcal{G}$ will factor as a product of potential functions that depend only on the cliques in $\mathcal{G}$; see for example Lauritzen [1996, ch. 3].

The above implies that inference of node-wise local conditional independences $Y_a \perp Y_{\mathcal{V}\setminus \mathrm{cl}(a)} | Y_{\mathrm{ne}(a)}$ suffices to identify the global conditional independences that hold in the joint distribution of $\mathbf{Y}$. These node-wise relationships will be learned through neighborhood selection [Meinshausen and Bühlmann, 2006].

3. Hurdle models. Univariate Hurdle models arise as modifications of a density through excision of points in the support, generally at the origin.

Suppose F is a measure on $\mathbb{R}$ admitting a density f with respect to some dominating measure λ. Then for any Borel set A, the **zero-modified** Hurdle measure F_0 on F is defined as

$$F_0(A) \equiv pF(A \setminus \{0\})/F(\{0\}^c) + (1-p)\delta_0(A),$$

where δ_0 is a point mass at 0. When λ is Lebesgue measure, as in the case in this work, $F(\{0\}) = 0$. This implies a density with respect to the sum-measure $\lambda + \delta_0$ of

$$f_0(y) = pf(y)v_y + (1-p)(1-v_y),$$

where $v_y = I_{y\neq 0}$ is the indicator function for non-zero values of y.

3.1. *Hurdle exponential families.* When F belongs to an exponential family, F_0 also belongs to an exponential family with modified sufficient statistics. In particular,when $f(y)$ is the Normal density with mean ξ and precision τ^2 then the Hurdle modification at zero has density

$$\begin{aligned} f_0(y) = \exp\big\{v_y \big[1/2\log\left(\tau^2/\left(2\pi\right)\right) + \log p/(1-p) - \xi^2\tau^2/2\big] \\ + y\xi\tau^2 - y^2\tau^2/2 + \log(1-p)\big\} \end{aligned} \tag{2}$$

with respect to the measure $\lambda + \delta_0$. This implies that f_0 is a member of an exponential family with sufficient statistics $v_y, y, -y^2/2$ and natural parameters

$$\begin{aligned} g &= 1/2\log\left(\tau^2/\left(2\pi\right)\right) + \log p/(1-p) - \xi^2\tau^2/2, \\ h &= \xi\tau^2, \\ k &= \tau^2, \end{aligned}$$

or inversely the original parameters in terms of the natural parameters are:

$$\begin{aligned} (\operatorname{Var} Y|v_y)^{-1} &= \tau^2 = k, \\ \operatorname{E} Y|v_y &= \xi = h/k, \\ \operatorname{logit}\operatorname{E} v_y &= \log p/(1-p) = g - 1/2\log\left(k/2\pi\right) + h^2/(2k). \end{aligned} \tag{3}$$

3.2. *Multivariate Hurdle models.* Based on figure 1, a plausible multivariate model puts positive mass on every one of the 2^m coordinate subspaces, including the origin and the entire Euclidean space $\mathbb{R}^m$. It is easiest to construct this model conditionally by first defining $\mathbf{V} = [V_1, \ldots, V_m]^T \equiv [I_{y_1\neq 0}, \ldots, I_{y_m\neq 0}]^T$ to be the vector of non-zero indicator functions on $\mathbf{Y}$, where the dependence on $\mathbf{Y}$ will be suppressed from here on. Under *any*

distribution on $\mathbf{Y}$, $\mathbf{V}$ is most generally a collection of Bernoulli variables distributed according to some 2^m probability table.

A random vector $\mathbf{Y}$ has singular Normal distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ [Rao, 1973] with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ with rank $r < m$ if $\mathbf{U}$ is a matrix such that $\mathbf{U}^T\boldsymbol{\Sigma} = 0$ and the following holds: a) $\mathbf{U}^T\mathbf{Y} = \mathbf{U}^T\boldsymbol{\mu}$ almost everywhere; and b) $\mathbf{Y}$ has a density $\frac{(2\pi)^{-r/2}}{(\det{}^+ \boldsymbol{\Sigma})^{1/2}} \exp\{-(\mathbf{y}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^-(\mathbf{y}-\boldsymbol{\mu})/2\}$, restricted to the hyperplane $\mathbf{U}^T\mathbf{Y} = \mathbf{U}^T\boldsymbol{\mu}$. Here $\det^+$ is the pseudo-determinant (product of non-zero eigenvalues) and $\boldsymbol{\Sigma}^-$ is a pseudo-inverse, such as the Moore-Penrose inverse. In the case that $\boldsymbol{\Sigma}$ is zero outside a positive-definite submatrix of size $r \times r$, $\mathbf{U}$ can be chosen to be a diagonal selection matrix consisting of zeros and ones, and $\mathbf{Y}$ has a density with respect to the measure $\lambda^r \otimes \delta_0^{m-r}$, which is the case treated subsequently.

In detail, given $\mathbf{V}$, let $r = |\mathbf{V}|$ and $\mathcal{I} \equiv \operatorname{diag}\mathbf{V}$ be the diagonal matrix with (i,i) entry equal to V_i, thus selecting the the non-zero entries of $\mathbf{V}$, then

$$(\mathbf{Y}|\mathbf{V} = \mathbf{v}) \sim \mathcal{N}\left(\mathcal{I}\boldsymbol{\mu}(\mathbf{v}), \mathcal{I}\boldsymbol{\Sigma}(\mathbf{v})\mathcal{I}\right),$$

where $\boldsymbol{\mu}(\mathbf{v})$ and $\boldsymbol{\Sigma}(\mathbf{v})$ depend on $\mathbf{v}$ arbitrarily, but are only identifiable along the obvious subspaces. This is the multivariate analog of excision of mass that defines the univariate Hurdle model. This construction is also equivalent to $\mathbf{Y}|\mathbf{V}$ being distributed as

$$(\mathbf{Y}|\mathbf{V} = \mathbf{v}) \sim \mathcal{N}\left(\mathbf{K}_\mathbf{v}^- \mathbf{h}_\mathbf{v}, \mathbf{K}_\mathbf{v}^-\right), \tag{4}$$

where $\mathbf{K}_\mathbf{v} = (\mathcal{I}\mathbf{K}(\mathbf{v})\mathcal{I})$ is a precision matrix $\mathbf{K}(\mathbf{v})$ with arbitrary dependence on $\mathbf{v}$ subject to the constraint that its rows and columns are filled with zeros whenever a coordinate of $\mathbf{v}$ is zero, $\mathbf{K}_\mathbf{v}^-$ is its pseudo-inverse, and $\mathbf{h}_\mathbf{v} = \mathcal{I}\mathbf{h}(\mathbf{v})$ is a column vector that may depend on $\mathbf{v}$ arbitrarily after being suitably zeroed out with its selection matrix.

The joint density $P(\mathbf{Y}, \mathbf{v}(\mathbf{Y})) = P(\mathbf{Y})$ can be decomposed as $P(\mathbf{Y}|\mathbf{V})P(\mathbf{V})$:

$$\begin{aligned}
p(\mathbf{y}) &= p(\mathbf{v}) \times (2\pi)^{-r/2}\left(\det{}^+ \mathbf{K}_\mathbf{v}\right)^{1/2} \exp\left\{-\left(\mathbf{y}-\mathbf{h}_\mathbf{v}\right)^T \mathbf{K}_\mathbf{v}\left(\mathbf{y}-\mathbf{h}_\mathbf{v}\right)/2\right\} \\
&= p(\mathbf{v}) \times (2\pi)^{-r/2}\left(\det{}^+ \mathbf{K}_\mathbf{v}\right)^{1/2} \exp\left\{-\mathbf{h}_\mathbf{v}^\mathbf{T}\mathbf{K}_\mathbf{v}\mathbf{h}_\mathbf{v}/2\right\} \\
&\quad \times \exp\left\{-\frac{1}{2}\mathbf{y}^T\mathbf{K}_\mathbf{v}\mathbf{y} + \mathbf{h}_\mathbf{v}^\mathbf{T}\mathbf{y}\right\} \\
&= \exp\left\{g(\mathbf{v}) - \mathbf{y}^T\mathbf{K}_\mathbf{v}\mathbf{y}/2 + \mathbf{h}_\mathbf{v}^\mathbf{T}\mathbf{y}\right\}, \qquad (5)
\end{aligned}$$

with respect to the m-fold product of $\lambda + \delta_0$.

As arbitrary functions of $\mathbf{v}$, $g(\mathbf{v})$, $\mathbf{h}_\mathbf{v}$ and $\mathbf{K}_\mathbf{v}$ are generically m-order polynomials in $\mathbf{v}$. When the polynomial order is 2, 1, 0 for $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{K}$

respectively, then the quadratic statistics $\mathbf{y}\mathbf{y}^T$, $\mathbf{v}\mathbf{y}^T$, $\mathbf{v}\mathbf{v}^T$ are minimally sufficient and the model simplifies to

$$p(\mathbf{y}) = \exp\left\{\mathbf{v}^T\mathbf{G}\mathbf{v} + \mathbf{v}^T\mathbf{H}\mathbf{y} - \frac{1}{2}\mathbf{y}^T\mathbf{K}\mathbf{y} - C\right\}, \tag{6}$$

where

$$C = \log \sum_{\mathbf{v}\in\{0,1\}^m} \exp\left[\mathbf{v}^T\mathbf{G}\mathbf{v} + \mathbf{h}^T\left(\mathcal{I}\mathbf{K}\mathcal{I}\right)^{-}\mathbf{h}/2\right]\left(\det{}^{+}\,\mathcal{I}\mathbf{K}\mathcal{I}\right)^{1/2}(2\pi)^{-r/2}.$$

It follows immediately that model (6) is a mixture of singular Gaussian distributions:

PROPOSITION 1. *Model (6) is a 2^m mixture of Normal distributions, each singular along a different coordinate axes. Given $\mathbf{v}$, the conditional expectation is $\left(\mathcal{I}\mathbf{K}\mathcal{I}\right)^{-}\mathbf{H}\mathbf{v}$ and the conditional variance is $\left(\mathcal{I}\mathbf{K}\mathcal{I}\right)^{-}$.*

3.3. *Conditional distributions identify interaction parameters.* The normalizing constant in equation (6) is a difficult to work with sum of 2^m terms, as is true also in the pairwise Ising model. However, the conditional likelihood of each coordinate $y_b|\mathbf{y}_A$ is tractable, and identifies parameters from a given row/column of the interaction matrices.

This is seen through examination of the conditional distributions. Recalling section 2.2, $\mathcal{V}$ denotes the vertex set. Fixing a coordinate b and its complement $A = \mathcal{V} \setminus b$, and noting that $v_i y_i = y_i$ and $v_i^2 = v_i$, the kernel of the distribution in (6) as a function of y_b is

$$\log f_{[b|A]}(y) = v_b g_{bb} + 2v_b\mathbf{g_{bA}}\mathbf{v}_A + v_b\mathbf{h_{bA}}\mathbf{y}_A + y_b h_{bb} + y_B\mathbf{h_{Ab}}^T\mathbf{v}_A - \frac{1}{2}\left[y_B^2 k_{bb} + 2y_b\mathbf{k_{bA}}\mathbf{y}_A\right] - C_{[b|A]}.$$

Factoring by the sufficient statistics over y_b,

$$\log f_{[b|A]}(y) = v_b\underbrace{\left[g_{bb} + 2\mathbf{g}_{bA}\mathbf{v}_A + h_{bA}\mathbf{y}_A\right]}_{g_{[b|A]}} + y_b^T\underbrace{\left[h_{bb} + \mathbf{h}_{Ab}^T\mathbf{v}_A - \mathbf{k}_{bA}\mathbf{y}_A\right]}_{h_{[b|A]}} - \frac{1}{2}y_B^2\underbrace{k_{bb}}_{k_{[b|A]}} - C_{[b|A]}. \tag{7}$$

Thus the conditional distribution is an example of the univariate Hurdle of equation (2) with natural parameters $g_{[b|A]}$, $h_{[b|A]}$ and $k_{[b|A]}$, which also

serve as linear predictors that depend on a design matrix constructed from $\mathbf{y}_A$ and $\mathbf{v}_A$.

Concretely, we can write

$$g_{[b|A]} = Z_0^T \theta_{g0} + \sum_{a \in A} X_a \theta_{ga},$$

where Z_0 in this case is taken to be 1, but generally could include a vector of covariates, $X_a = [v_a, y_a]$ and $\theta_{ga} = [g_{ba}, h_{ba}]$. The linear predictor for $h_{[b|A]}$ can be written analogously. We can solve for the natural parameters to yield parameters recognizable from the univariate case using equation (3). Then

$$\begin{aligned}
\tau^2_{b|A]} &= k_{bb}, \\
\xi_{[b|A]} &= \frac{1}{k_{bb}} \left[h_{bb} - \mathbf{k_{Ab}}^T \mathbf{y}_A + \mathbf{h_{Ab}}^T \mathbf{v}_A \right], \\
[\log p/(1-p)]_{[b|A]} &= [g_{bb} + 2\mathbf{g_{bA}} v_a + \mathbf{h_{bA}} \mathbf{y}_A] \\
&\quad \underbrace{-1/2 \log\left(k_{bb}/2\pi\right) + \left[h_{bb} - \mathbf{k_{Ab}}^T \mathbf{y}_A + \mathbf{h_{Ab}}^T \mathbf{v}_A \right]^2 / \left(2 k_{bb}\right)}_{C_{[b|A]}},
\end{aligned}$$

so $\log p/(1-p)$ depends on y_A both through a linear term $2g_{bA}\mathbf{v}_A + h_{bA}\mathbf{y}_A$ as well as a quadratic term (derived from the normalizing constant for the Gaussian). Integrating with respect to $\delta_0 + \lambda$ confirms directly that the normalizing constant is indeed given by $C_{[b|A]}$.

The conditional distribution in (7) defines a vector generalized linear model, since the univariate family (2) is parametrized by three natural parameters, g, h and k, the first two of which are modeled as a linear function of covariates.

3.4. *Other work on mixed graphical models.* The notation used here follows Lauritzen [1996], who describes conditional Gaussian (CG) models with *inhomogeneous, non-singular* precision $\mathbf{K}(\mathbf{v})$ that can depend on the discrete set of covariates in arbitrary, positive-definite fashion. The formulation here is both a special case, and an extension of Lauritzen's inhomogeneous CG models, since it imposes the inhomogeneity (conditional singularity, in fact) in a structured fashion.

Several authors have described algorithms to infer the structure of specializations of Lauritzen's CG models. Lee and Hastie [2013] and Cheng et al. [2013] describe ℓ_1-penalized, pseudo-likelihood algorithms to estimate structure in specializations of the inhomogeneous, pairwise case (in which the variance of continuous variables, given both discrete and continuous neighbors are homoscedastic).

Other authors have described classes of graphical models with mixed node conditional distributions. Chen et al. [2015], and Yang et al. [2014] describe exponential-family graphical models for which the conditional distributions follow a generalized linear model with neighbors entering additively. While this paper was in preparation, Tansey et al. [2015] proposed *vector space graphical models* that include the multivariate Hurdle as a special case, estimated through sparse group-lasso penalized neighborhood selection. The isometric group-lasso does not account for heterogeneity in the scaling of predictors in the conditional distributions. The anisometric group-lasso proposed in the following section yields substantial benefits in finite samples.

Other work considering dependence measures in mixed distributions include Olkin and Tate [1961], Cox and Wermuth [1992].

4. Neighborhood estimation via penalized regression. Equation (7) shows that the conditional distributions of each node, given the rest, identify rows and columns of the interaction matrices. If m is fixed and $n \to \infty$, maximum likelihood, methods of moments or Bayesian estimators will concentrate around the true values and hypothesis testing might reveal the neighborhood of a node. Although in single cell experiments, the number of cell replicates, n, is larger than in many bulk mRNA experiments, it is still generally the case that the number of genes, m measures in the thousands, while n measures in the hundreds (though emerging technologies may change this).

On the other hand, under scenarios in which $n, m \to \infty$ while $n > Cd^{\phi}(\log m)^{\psi}$, where C, ϕ and ψ are constants that depend on the model and d is the maximum vertex degree, penalized regression has been shown to consistently identify signed sparsity patterns in precision matrices in multivariate Normal models [Meinshausen and Bühlmann, 2006], in interaction matrices for Ising (auto-logistic) graphical models [Ravikumar et al., 2010] and exponential family graphical models [Yang et al., 2014]. This motivates the application of node-wise penalized regression for the problem at hand.

4.1. *Penalty and computation.* By inspection of (7) for $y_b \perp y_a | y_{\mathcal{V}\setminus\{a,b\}}$ the four parameters $[g_{ba}, h_{ba}, h_{ab}, k_{ba}] = \theta_a$ must simultaneously vanish. Penalizing the conditional log-likelihood $\log f_{[b|A]}(\mathbf{y})$ with the grouped ℓ_1 penalty $P_\lambda(\theta) = \lambda \sum_{a \in A} \sqrt{\theta_a^T \theta_a}$ can lead to a solution that is sparse in parameter blocks responsible for vertex a. This penalty is equivalent to placing a sequence of independent, multivariate Laplace (multivariate exponential power distribution [Eltoft et al., 2006]) priors on blocks of θ and reporting the MAP. It is well-known that this results in both shrinkage and variable selection.

Viewed as a prior, the default group-lasso penalty implicitly assumes that each variable in each block has a similar effect size. This may be reasonable, provided they are measured in comparable units. For example if covariate X_1 is measured in meters, while covariate X_2 in centimeters, then the distribution of effect sizes for X_2 would be 1000-times more dispersed than the distribution for X_1, revealing a kind of expected scale-equivariance. In penalized GLMs, this is typically enforced "at run time" by ensuring covariates are on comparable scales, or $Z-$scoring each column of the design if no intrinsic scale exists.

In the case of a vector regression, terms from linear predictor $g_{[b|A]}$ and linear predictor $h_{[b|A]}$ end up together in blocks, and these coefficients are not necessarily comparable, as one specifies log-odds of $E(V_b|V_a = 0)$ while the other specifies conditional expectations of $E(Y_a|Y_b)$. Re-scaling is not an option, since the same design matrix $X_a = [V_a, Y_a]$ is used in each linear predictor, and even if it were, only reparametrization through isometric transformations produces the same solution (in terms of fitted values) as has been noted for the group-lasso in linear regression [Simon and Tibshirani, 2012]. But replacing the isometric ℓ_2 norm in the sum so that the penalty is

$$P_{\mathbf{H},\lambda}(\theta) = \lambda \sum_{a \in A} \sqrt{\theta_a^T \mathbf{H}_{aa} \theta_a}, \tag{8}$$

where $\mathbf{H} \equiv \text{diag}\,(\mathbf{H}_{aa})$ is a block-diagonal, positive-definite matrix allows terms from the linear predictors to have different scales of penalty. It also accounts for correlation between components of θ_a, since columns of the design are correlated due to both v_a and y_a appearing as predictors.

If prior information existed, the matrix $\mathbf{H}$ could be chosen accordingly, with interpretation as a multivariate Laplace prior. Absent prior information, setting $\mathbf{H}$ equal to the Fisher information under a null model $\theta_a = 0$ for all a results in variable selection approximately equal to conducting score tests, with exact equivalence holding under a null hypothesis of $\theta_a = 0$ for all a, as is shown in the following proposition:

PROPOSITION 2. *Let* $\mathbf{H} = \left[\frac{\partial^2 \log f_{[b|A]}(\mathbf{y})}{\partial \theta_i \theta_j}\right]$ *be the conditional information. Suppose that the inverse information* $\mathbf{H}^{-1}$ *is block-diagonal. Then the scaled* ℓ_1 *penalty is equivalent to a score test of the null hypothesis that* $\theta = 0$ *vs the alternative that a pre-specified subvector* $\theta_a \neq 0$.

Proof: Let $c = \mathcal{V} \setminus \{a, b\}$ and suppose that $\theta_c = 0$. From the KKT conditions, $\theta_a = 0$ is an optimum if and only if

$$\nabla_a^T H_{aa}^{-1} \nabla_a < \lambda^2,$$

where $\nabla_a = \frac{\partial \log f_{[b|A]}(\mathbf{y})}{\partial \theta_a}$ is the a-subvector of the conditional log-likelihood gradient. Taking λ^2 to be an appropriate quantile from a χ^2-distribution with $\dim(H_{aa})$ degrees of freedom results yields a score test.

This leads to algorithm 1. The covariates $\mathbf{Z}$ might just be an intercept column, but generally could be any cell-level covariate deemed relevant, in which case the estimated model is a *conditional Markov random field*. The smooth and concave function in line 3 can be maximized using any Newton-like algorithm (eg BFGS). The objective in line 6 is a sum of a concave, smooth function and a structured concave function and can be efficiently solved using proximal gradient ascent [Parikh and Boyd, 2014]. In particular, one may exploit the fact that although the proximal operator

$$\text{prox}_\gamma(x) = \text{argmax}_u \frac{1}{\gamma}\|x-u\|_2^2 + \sum_{a\in A}\sqrt{u_a^T H_{aa} u_a}$$

is not available in the familiar form of a soft-thresholding operator as in the isometric group-lasso, the proximal operator of the *anisometric* group-lasso can be efficiently found via a line search and a few 4×4 matrix multiplications after one-time pre-calculation of the singular value decomposition of H_{aa} [Foygel and Drton, 2010]. Throughout the inner-loop, warm starts are exploited for $\hat{\theta}$ as λ varies. Active set heuristics using the strong rules of Tibshirani et al. [2012] yield computational gains for sparse solutions with large m. In the accompanying software, the algorithm is written in a combination of `R` and `C++`.

Data: Expression matrix $\mathbf{Y} \in \mathbb{R}^{n\times m}$, covariates $\mathbf{Z} \in \mathbb{R}^{n\times q}$, penalty path $\mathbf{\Lambda}$
Result: Neighborhoods $ne(i,\lambda)$, $1 \le i \le m$, $\lambda \in \mathbf{\Lambda}$
1 **for** $i = 1 \dots m$ **do**
2 $\mathbf{X} \leftarrow \left[\mathbf{Z}, V_1, Y_1, \dots, V_{i-1}, Y_{i-1}, V_{i+1}, Y_{i+1}, \dots, V_m, Y_m\right]$;
3 $\bar{\theta}_0 \leftarrow \text{argmax}_{\theta_0, \theta=0} \log f([\theta_0, \theta], \mathbf{X})$;
4 $\mathbf{H} \leftarrow \text{Hessian}\big(\log f\left([\bar{\theta}_0, 0], \mathbf{X}\right)\big)$;
5 **for** $\lambda \in \mathbf{\Lambda}$ **do**
6 $[\hat{\theta}_0, \hat{\theta}] \leftarrow \text{argmax}_{\theta_0,\theta} \log f([\theta_0,\theta],\mathbf{X}) - P_{\lambda,\mathbf{H}}(\theta)$;
7 $ne(i,\lambda) \leftarrow \text{ProcessNeighborhood}(\hat{\theta})$;
8 **end**
9 **end**

Algorithm 1: Neighborhood selection

4.2. *Convergence and model selection consistency.* Tansey et al. [2015] provide model selection consistency guarantees under assumptions involving both the sample information matrix and the joint log-partition function. The

log-partition function in model (6) is computationally intractable for even moderate m thus these assumptions are generally difficult to verify, even in simulation. In practice, total selection consistency may require astronomically large samples; for example in the scenario described in section 5, it fails to occur even once in 30 simulations for $m = 128, n = 10000$. Empirical selection consistency also does not hold in the large sample ($n = 25000$) simulations of Tansey et al. [2015].

However, at realistic sample sizes imperfect recovery, in which some number of false edges are included in a portion of the true set, is feasible, and is explored further in section 5. Applying a positive definite penalty matrix rather than the isometric penalty offers drastically improved rates of imperfect recovery.

5. Simulations. We consider a series of simulations under two sets of parametric alternatives. Observations are generated through Gibbs sampling from model (6), with 2000 iteration burn-in, 10% down-sampling. Down-sampled iterations exhibited only mild auto-correlation.

In the first study, the dependence structure is blockwise-dense: if any of $g_{ij}, h_{ij}, h_{ji}, k_{ij} \neq 0$, then all of them are non-zero. In the second study, only g_{ij} is non-zero, and the Hurdle model is a superset of the true model, which is Ising/logistic. In both cases the edge density is fixed at 1.5% and the underlying graph is a chain, thus the maximum degree is 2. The number of observations $n = 100$ is fixed and the dimension varies from $m = 16$ to $m = 128$, with 30 replicates run.

Five methods were examined to test graph structure inference: neighborhood selection under the Hurdle model using 1) isometric and 2) anisometric penalties, neighborhood selection using ℓ_1-penalized 3) logistic regression, 4) linear regression and 5) Gaussian copula (NPN) models [Liu et al., 2009]. The logistic model is fit using the R package `glmnet`. Models 4 and 5 were fit using the R package `huge`, with "truncation" transformation function. Neighborhoods are stitched together using an "or" rule, i.e. vertices a and b are adjacent if either $b \in \text{ne}(a)$ or $a \in \text{ne}(b)$.

In each method, as a tuning parameter decreases, more edges are selected. Figure 2 shows a representative ROC curve (true positives vs false positives) for $n = 100$ and $m = 64$. In figure 3, the maximum sensitivity $\left(\frac{\text{true positives}}{\text{total true}}\right)$ under the oracle value of the tuning parameter that admits fewer than 10% false discoveries $\left(\frac{\text{false positives}}{\text{total positives}}\right)$ is shown.

These simulations show that mis-specified estimation procedures perform poorly when model (6) is the data generating distribution. When a subset of model (6) holds, as in the sparse dependence scenario, the over-parametrized

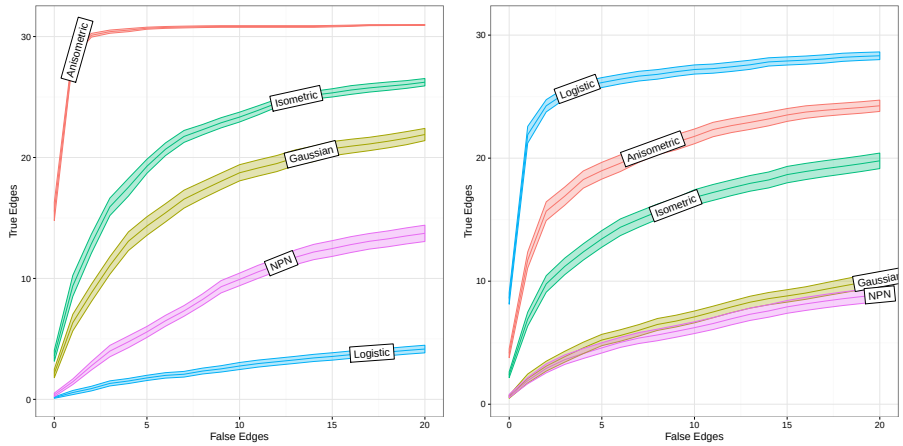

Fig 2: Number of true positives vs false positives for simulated chain graphs under *dense* and *sparse dependence* with $m = 64$ nodes and $n = 100$ observations. The ribbon shows the simulation-induced standard errors about the average. The Anisometric and Isometric models use neighborhood selection with the multivariate Hurdle model (7) with group-ℓ_1 penalty based on the null-model Fisher information and identity matrix, respectively. The Gaussian, NPN and Logistic models use ℓ_1 penalized neighborhood selection under linear (Gaussian), Normal-score transformed linear (NPN) and logistic regressions.

Hurdle model recovers fewer edges than the minimal Logstic model, but the anisometric ℓ_1 penalty partially closes this gap. The simulations also suggest that perfect recovery of gene networks is impractical at realistic sample sizes, even with a correctly specified model, motivating a meta-analysis of estimated graphs, which is proposed in section 7.2.

6. T follicular helper cells. We consider co-expression networks in Tfh cells measured in eight healthy donors. 65 genes were selected for profiling via qPCR on the basis of their role in Tfh signaling and differentiation, generally with sparse expression across single cells (median expression 18%, 90th percentile, 71%). 465 single cell, and 187 10-cell replicates were taken.

Figure 4 shows networks of 22 edges estimated using Hurdle, Gaussian (with centered data, see section 9) and Logistic, and Gaussian model using 10-cell aggregates. The size of the network is a compromise between stability selected [Shah and Samworth, 2013] sizes of each procedure, which vary from 11 (Hurdle), 12 (Gaussian10), 14 (Logistic) to 32 edges (Gaussian).

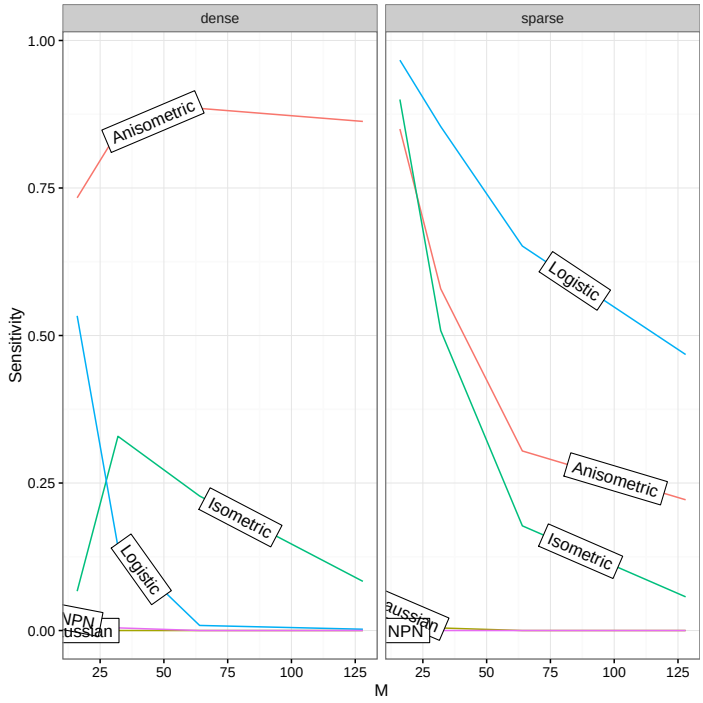

Fig 3: Sensitivities $\left(\frac{\text{true positives}}{\text{total true}}\right)$ at 10% FDR as m, the number of nodes increases for a chain graph under fixed sparsity. See the caption of figure 2 for a description of the methods.

	Gaussian(10)	Hurdle	logistic	Gaussian(raw)
Gaussian	1.00	**0.69**	1.00	0.96
Gaussian(10)	-	0.84	1.00	1.00
Hurdle	-	-	1.00	1.00
logistic	-	-	-	**0.28**

TABLE 1

Dissimilarities $\left(\frac{Hamming\ Distance}{Number\ of\ edges}\right)$ *between networks of size 22 estimated through various methods. The Gaussian(10) model is a Gaussian model estimated on 10-cell replicates, while the Gaussian(raw) data is estimated on single cells without centering the data. The Hurdle and logistic models are described in the text.*

Normalized Hamming distances between the four methods and the Gaussian model fit on the "raw", uncentered data are reported in table 1. The Hurdle and (centered) Gaussian models are most similar, while the logistic and Gaussian 10-cell network are quite distinct. The Gaussian(raw) model on untransformed data is similar to the logistic model, as distance of non-zero expression values from the origin is large compared to the variation among the non-zero values.

In the Hurdle network, the transcription factors NFATC1 (Nuclear factor of activated T-cells) and BCL6, and the signaling molecule CD154 and chemokine receptor CCR3 are hubs. NFATC1 has been found to promote transcription of cytokines IL21 [Hermann-Kleiter and Baier, 2010] and signaling molecule CD154 [Pham et al., 2005], while BCL6 serves as a transcriptional repressor, and is one of the canonical markers constitutively expressed in Tfh cells. CTLA4 which has been described to inhibit inflammation, interacts negatively with inflammatory activator JAK3. The disconnected component of CCR3-CCR4-BTLA-SELL-TNFSF4 may hint at plasticity between Tfh cells and the related T-cell lineages Th1 and Th2. CCR3 and CCR4 are canonical markers of Th2 cells, while TNFSF4 (coding for OX40L) promotes Th2 development de Jong et al. [2002]. Thus co-expression of these genes may suggest cells transitioning between Tfh and Th1 or Th2 states.

In the Gaussian network, though NFATC1, BCL6 and CD154 remain highly connected, CD27 now has highest degree and serves as a hub to receptors CXCR4, IL2Rb, IL2Rg, as well as ITGB2, NFATC1 and FYN. CD3e, the backbone responsible for transducing the T-cell receptor signal is connected with co-receptor CD4, CD154, IL2Rg, Fyn and ANP32B. The negative interactions between BTLA and CTLA4 are absent.

The logistic network consists primarily of negative interactions. The strongly negative BCL6–BLIMP1 edge is consistent with previously described antagonism between these genes [Johnston et al., 2009]. Interestingly, this edge is absent in the other networks.

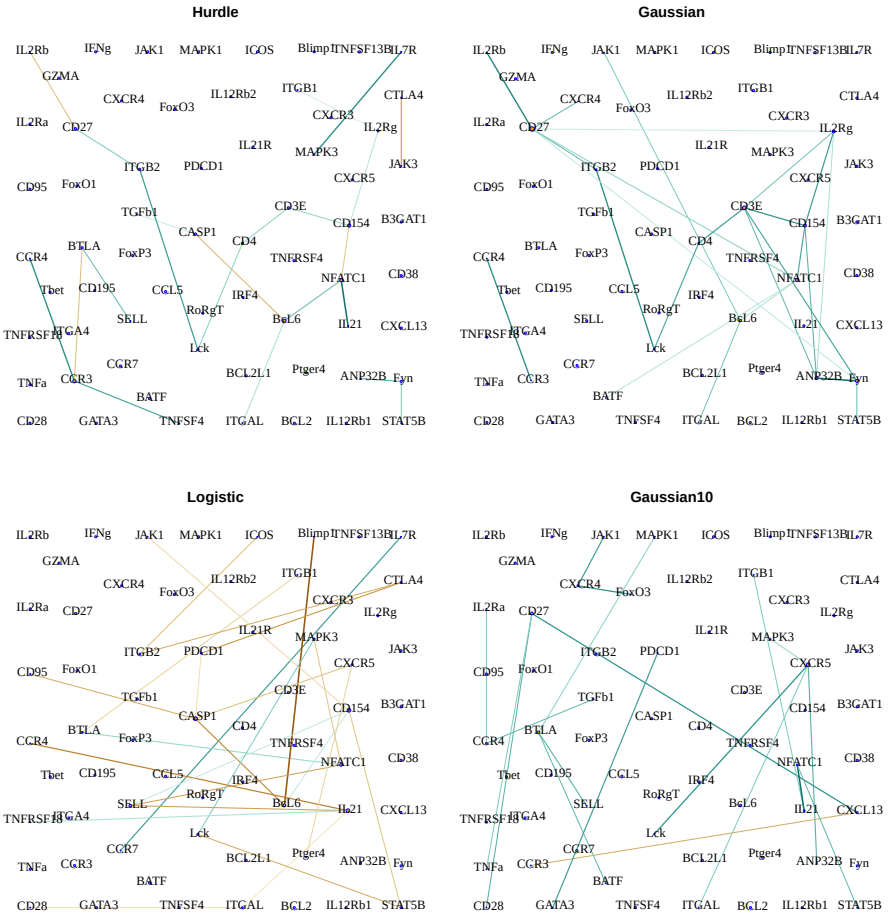

Fig 4: Networks of 22 edges estimated through neighborhood selection under the Hurdle, logistic, Gaussian model (single cells) and Gaussian model (10 cell aggregates) in T follicular helper cells. Brown hues indicate estimated negative dependences, while blue-green hues indicate positive dependences. The edge width and saturation are larger for stronger estimated dependences.

7. Mouse dendritic cells. In an experiment originally described in Shalek et al. [2014], bone marrow-derived dendritic cells, from *mus musculus* were exposed to lipopolysaccharide (LPS), a toxic compound secreted and structurally utilized by gram-negative bacteria. Cells were sampled after 0, 1, 2, 4, and 6 hours. Here we consider the transcription networks estimated using 4431 transcripts expressed in at least 20% of 65 cells sampled 2 hours after LPS exposure. Rather than attempting to perform model selection on this limited sample size, we consider highly sparse ($< .01\%$ sparsity) networks of 700 edges, chosen to provide tractable visualization and illustration of the method.

7.1. *Selected Networks.* In a Gaussian model, the network is star-shaped, with Mx1, Ccl17, Tax1bp3 and Ccl3 as hubs all with degrees ≥ 15, though none are directly inter-connected (figure 5). In all, 2.5% of non-isolated vertices contribute 50% of the edges in the network. With the exception of Tax1bp3, these hub genes are all immune-signaling related.

In the Hurdle model (figure 6), the graph is chain-shaped, with the maximum degree being 12: 7% of nodes provide 50% of the edges. The strongest hub, Mgl2 (also known as Cd301b), has been recently described to be involved in uptake and presentation of glycosylated antigens, such as LPS, by dendritic cells [Denda-Nagai et al., 2010]. A sub-connected set of genes coding for MHC-II antigen presentation (H2ab1, H2eb1, H2aa) is the densest sub-component, and interconnected to Mgl2 as well as Fabp5. Increased expression of Fabp5 has been shown to increase expression of cytokines Il7 and Il18, hence is also involved in immune cell stimulation [Adachi et al., 2012]. Many of the neighbors of Mgl2,H2ab1, H2eb1, H2aa and Fabp5 are neighbors of the hub genes in the Gaussian graph, whereas Mx1, Ccl17 and Ccl3 are sparsely connected in the Hurdle network. Tax1bp3 is absent.

7.2. *Graphical geneset edge enrichment.* We consider how well the 700 edges recapitulate known relationships between genes using the Gene Ontology (GO) annotations. The Gene Ontology Consortium [2015] provides a directed, acyclic graph of ontologies to which genes may be annotated if they have been shown experimentally or computationally to be involved in a biological process, component or function. In the GO annotations for mouse, most genes are members of several categories (median= 14). We test for enrichment between and within categories under a hypergeometric model, in which each pair (i, j) of (possibly non-disjoint) GO categories induces a *coloring* $(i, j) \to c$ of vertices, coloring any vertex belonging to either i or j. Evidence that this color is interconnected, given 700 total edges, n_c of which are between c-colored vertices, is evaluated using the hypergeometric

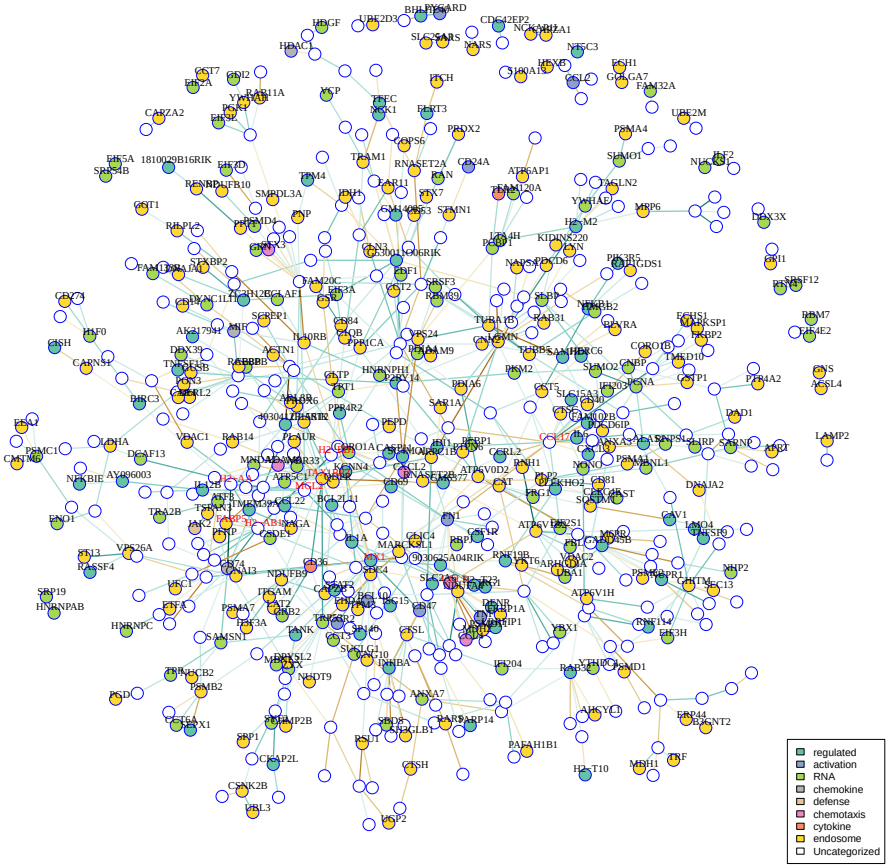

Fig 5: Gaussian model networks in LPS-treated mouse dendritic cells. Hub genes are shown in red. Vertex colors indicate gene ontology membership.

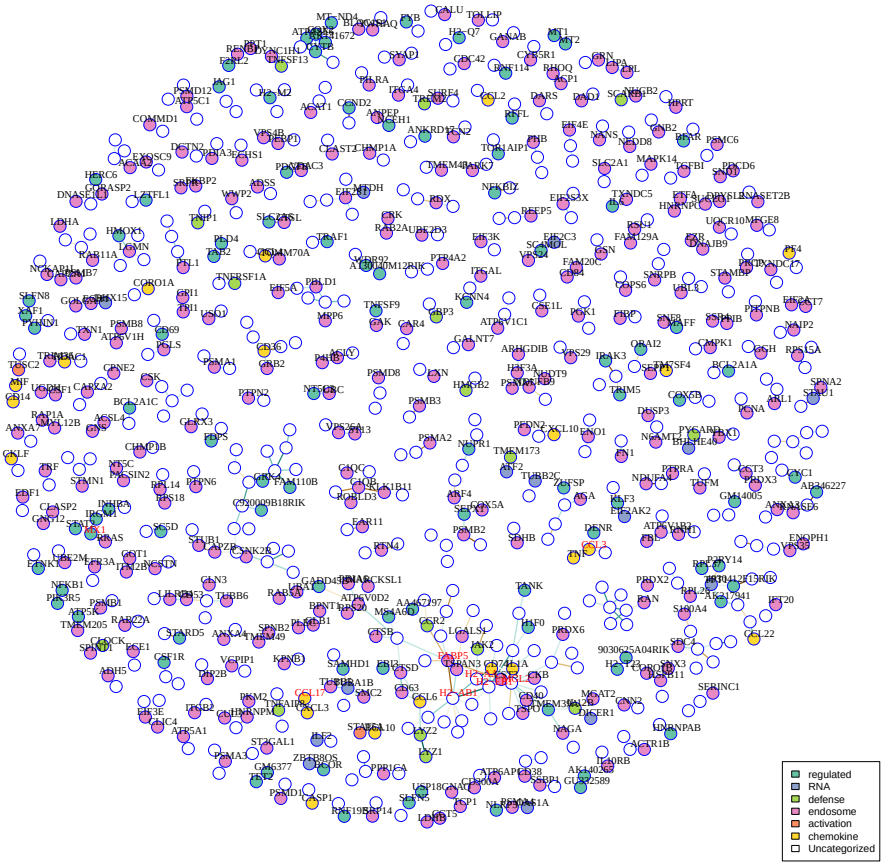

Fig 6: Hurdle model networks estimated in LPS-treated mouse dendritic cells. Hub genes are shown in red. Vertex colors indicate gene ontology membership.

tail probability $t(c) = P(N_c > n_n)$ in an urn of $4431^2/2$ possible edges, of which m_c could possibly join c-colored vertices. The distribution of the smallest $k < 200$ order statistics $t_{(k)}$ across the ~ 16 million dependent tests possible on $3987^2/2$ pairs of categories is found under a Erdos-Renyi random graph model, yielding Monte Carlo p-values $p(t_{(k)})$. A color is declared significant if $p(t_c) < .05$ and all colors ranked above it are also significant at .05 (using the ordering on $t(c)$), thus controlling the FDR at less than 5%.

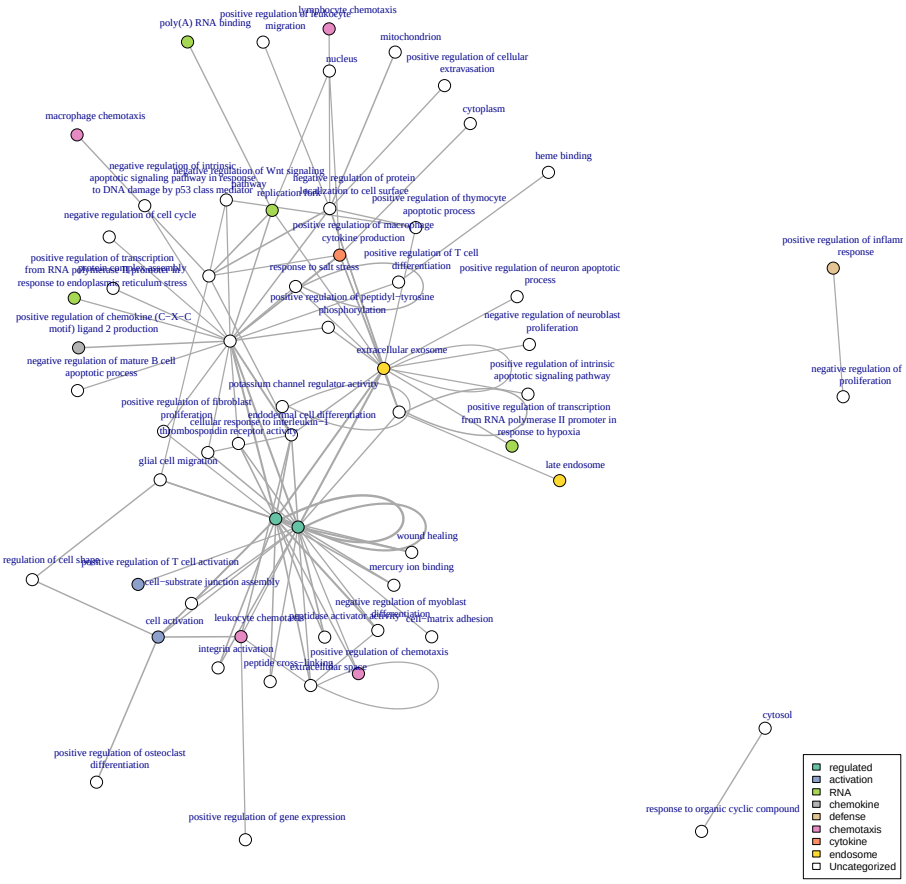

Fig 7: Modules enriched at FDR $\leq 5\%$ using graphical geneset edge enrichment in mouse dendritic cells under Gaussian model.

In the Gaussian model, more than 100 pairs of categories (colors) are significantly enriched at an FDR of less than 5%, however in these 100

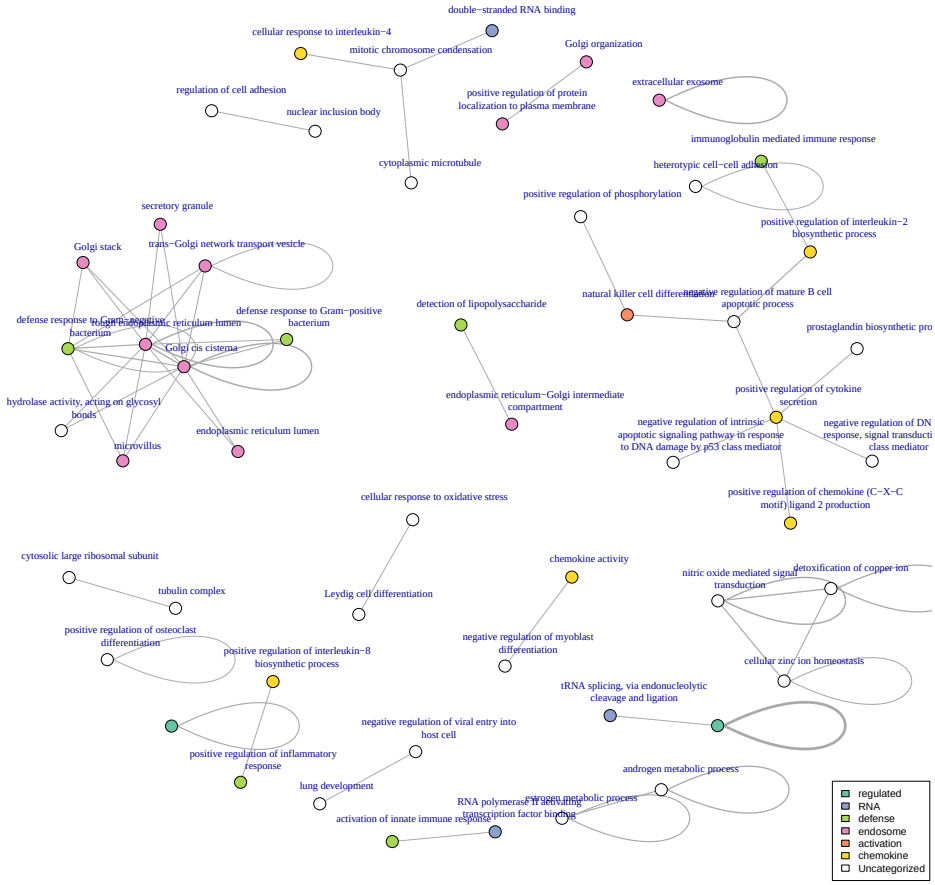
double−stranded RNA binding
cellular response to interleukin−4
mitotic chromosome condensation
Golgi organization
regulation of cell adhesion
nuclear inclusion body
positive regulation of protein localization to plasma membrane
extracellular exosome
immunoglobulin mediated immune response
cytoplasmic microtubule
heterotypic cell−cell adhesion
positive regulation of phosphorylation
secretory granule
positive regulation of interleukin−2 biosynthetic process
Golgi stack
trans−Golgi network transport vesicle
natural killer cell differentiation
negative regulation of mature B cell apoptotic process
defense response to Gram−negative bacterium
rough endoplasmic reticulum lumen
defense response to Gram−positive bacterium
detection of lipopolysaccharide
prostaglandin biosynthetic pro
Golgi cis cisterna
endoplasmic reticulum−Golgi intermediate compartment
positive regulation of cytokine secretion
hydrolase activity, acting on glycosyl bonds
negative regulation of intrinsic apoptotic signaling pathway in response to DNA damage by p53 class mediator
negative regulation of DN response, signal transducti class mediator
endoplasmic reticulum lumen
microvillus
positive regulation of chemokine (C−X−C motif) ligand 2 production
cellular response to oxidative stress
cytosolic large ribosomal subunit
chemokine activity
detoxification of copper ion
nitric oxide mediated signal transduction
tubulin complex
Leydig cell differentiation
positive regulation of osteoclast differentiation
negative regulation of myoblast differentiation
positive regulation of interleukin−8 biosynthetic process
cellular zinc ion homeostasis
tRNA splicing, via endonucleolytic cleavage and ligation
negative regulation of viral entry into host cell
positive regulation of inflammatory response
androgen metabolic process
lung development
estrogen metabolic process
RNA polymerase II activating transcription factor binding
activation of innate immune response
regulated
RNA
defense
endosome
activation
chemokine
Uncategorized

Fig 8: Modules enriched at FDR $\leq 5\%$ using graphical geneset edge enrichment in mouse dendritic cells under Hurdle model.

pairs, only 7 correspond to intra-category enrichment (figure 7). These are: response to salt stress, potassium channel regulator activity, extracellular space, extracellular exosome and three manually curated genesets containing genes with significant time-course differential expression in the original experiment. In the Hurdle model (figure 8), 14/60 significantly enriched pairs form intra-connections, including defense response to Gram-negative bacteria, and cell-cell adhesion and several modules involving extracellular secretion via the Golgi apparatus. Also of particular note, genes annotated to the activation of innate immune response are directly connected to RNA PolII transcription factors, a connection absent in the Gaussian model. This suggests that using the more parametrically appropriate Hurdle model manages to identify transcription factor-induced expression changes in these regulated genes, a direct method by which one gene would induce expression changes in another. No significant enrichment was found in the logistic model.

8. Discussion. Graphical models estimated from single cell data are distinct from networks estimated from bulk data, or even repeated stochastic samples. In simulations, the Hurdle model with anisometric penalty has much greater sensitivity compared to available methods, while in the two data sets here, it yields substantially different network estimates compared to Gaussian and Logistic models on these zero-inflated data. When enrichment of gene ontology categories is considered between vertices in transcriptome-wide data, the enrichment uncovered with the Hurdle model is consistent with identifying direct effects of transcription factors on genes undergoing dynamic regulation due to LPS exposure.

Although measuring transcriptome-wide data allows conditional estimation of direct effects between genes, non-mRNA factors may also greatly affect gene expression. In this sense, important variables have still been marginalized over, and in the case of the Tfh data, indeed, most of the transcriptome has been marginalized over. Thus, co-expression under sparsity assumptions is most helpful as a screening procedure. Methods to adapt graphical model variable selection to clustering and/or factor analytic models would extend its usefulness greatly, and allow greater biological insight with these data sets.

9. Supplemental methods. In all models and data sets, the cellular detection rate $\sum_j I_{y_{ij}>0}$ [Finak et al., 2015] was used as an unpenalized covariate in $\mathbf{Z}$ as described in algorithm 1. In the Tfh data, a separate intercept was fit for donor, as well. For the Gaussian and Hurdle models,

positive values were conditionally centered

$$\tilde{y}_{ij} = \begin{cases} 0 & v_{ij} = 0, \\ y_{ij} - \bar{y}_j^+ & \text{else}, \end{cases}$$

where $\bar{y}_j^+$ is the average in a gene over positive values, to make V_j and Y_j marginally orthogonal, which enhanced optimization convergence, and reduced the leverage of zeros in the Gaussian model.

The mDC data set was thresholded as described previously [Finak et al., 2015], and filtered for low-expression and cluster-disrupted cells.

Acknowledgments. AM and RG acknowledge funding through grant R01 EB008400 from the National Institute of Biomedical Imaging and Bioengineering, US National Institutes of Health.

MD was partially supported by grant DMS 1561814 from the US National Science Foundation.

AM thanks Daniel Lu for comments on the networks in reported section 6.

References.

Yasuhiro Adachi, Sumie Hiramatsu, Nobuko Tokuda, Kazem Sharifi, Majid Ebrahimi, Ariful Islam, Yoshiteru Kagawa, Linda Koshy Vaidyan, Tomoo Sawada, Kimikazu Hamano, and Yuji Owada. Fatty acid-binding protein 4 (FABP4) and FABP5 modulate cytokine production in the mouse thymic epithelial cells. *Histochemistry and Cell Biology*, 138 (3):397–406, 2012.

Barry Arnold and James Press. Compatible conditional distributions. *Journal of the American Statistical Association*, 84(405), 1989.

Shizhe Chen, Daniela M. Witten, and Ali Shojaie. Selection and estimation for mixed graphical models. *Biometrika*, 102(1):47–64, nov 2015.

Jie Cheng, Elizaveta Levina, and Ji Zhu. High-dimensional mixed graphical models. *arXiv preprint arXiv:1304.2810*, apr 2013.

D.R. Cox and Nanny Wermuth. Response models for mixed binary and quantitative variables. *Biometrika*, 79(3):441–461, 1992.

Esther C de Jong, Pedro L Vieira, Pawel Kalinski, Joost H N Schuitemaker, Yuetsu Tanaka, Eddy a Wierenga, Maria Yazdanbakhsh, and Martien L Kapsenberg. Microbial compounds selectively induce Th1 cell-promoting or Th2 cell-promoting dendritic cells in vitro with diverse th cell-polarizing signals. *Journal of Immunology*, 168(4): 1704–1709, 2002.

Kaori Denda-Nagai, Satoshi Aida, Kengo Saba, Kiwamu Suzuki, Saya Moriyama, Sarawut Oo-puthinan, Makoto Tsuiji, Akiko Morikawa, Yosuke Kumamoto, Daisuke Sugiura, Akihiko Kudo, Yoshihiro Akimoto, Hayato Kawakami, Nicolai V. Bovin, and Tatsuro Irimura. Distribution and function of macrophage galactose-type C-type lectin 2 (MGL2/CD301b): Efficient uptake and presentation of glycosylated antigens by dendritic cells. *Journal of Biological Chemistry*, 285(25):19193–19204, 2010.

Adrian Dobra, Chris Hans, Beatrix Jones, Joseph R. Nevins, Guang Yao, and Mike West. Sparse graphical models for exploring gene expression data. *Journal of Multivariate Analysis*, 90:196–212, 2004.

Mathias Drton and Marloes Maathuis. Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, (in press), 2017.

Torbjørn Eltoft, Taesu Kim, and Te Won Lee. On the multivariate Laplace distribution. *IEEE Signal Processing Letters*, 13(5):300–303, 2006.

Greg Finak, Andrew McDavid, Masanao Yajima, Jingyuan Deng, Vivian Gersuk, Alex K. Shalek, Chloe K. Slichter, Hannah W. Miller, M. Juliana McElrath, Martin Prlic, Peter S. Linsley, and Raphael Gottardo. MAST: a flexible statistical framework for assessing transcriptional changes and characterizing heterogeneity in single-cell RNA sequencing data. *Genome Biology*, 16(1):278, 2015.

Rina Foygel and Mathias Drton. Exact block-wise optimization in group lasso and sparse group lasso for linear regression. *Arxiv preprint arXiv:1010.3320*, pages 1–19, 2010.

Natascha Hermann-Kleiter and Gottfried Baier. NFAT pulls the strings during CD4+ T helper cell effector functions, 2010.

Kevin a Janes, Chun-Chao Wang, Karin J Holmberg, Kristin Cabral, and Joan S Brugge. Identifying single-cell molecular programs by stochastic profiling. *Nature Methods*, 7 (4):311–317, 2010.

Robert J Johnston, Amanda C Poholek, Daniel DiToro, Isharat Yusuf, Danelle Eto, Burton Barnett, Alexander L Dent, Joe Craft, and Shane Crotty. Bcl6 and Blimp-1 are reciprocal and antagonistic regulators of T follicular helper cell differentiation. *Science*, 325, 2009.

Jong Kyoung Kim and John C Marioni. Inferring the kinetics of stochastic gene expression from single-cell RNA-sequencing data. *Genome Biology*, 14(1), 2013.

Steffen Lauritzen. *Graphical Models*. Oxford University Press, Oxford, 1st edition, 1996.

Jason D Lee and Trevor J Hastie. Structure Learning of Mixed Graphical Models. In *AISTATS 16*, volume 31, pages 388–396, Scottsdale, AZ, USA, 2013.

Lin Lin, Greg Finak, Kevin Ushey, Chetan Seshadri, Thomas R Hawn, Nicole Frahm, Thomas J Scriba, Hassan Mahomed, Willem Hanekom, Pierre-Alexandre Bart, Giuseppe Pantaleo, Georgia D Tomaras, Supachai Rerks-Ngarm, Jaranit Kaewkungwal, Sorachai Nitayaphan, Punnee Pitisuttithum, Nelson L Michael, Jerome H Kim, Merlin L Robb, Robert J O'Connell, Nicos Karasavvas, Peter Gilbert, Stephen C De Rosa, M Juliana McElrath, and Raphael Gottardo. COMPASS identifies T-cell subsets correlated with clinical outcomes. *Nature Biotechnology*, 33(6):610–616, 2015.

Han Liu, John Lafferty, and Larry Wasserman. The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs. *Journal of Machine Learning Research*, 10(10):2295–2328, 2009.

F. J. Livesey. Strategies for microarray analysis of limiting amounts of RNA. *Briefings in Functional Genomics and Proteomics*, 2(1):31–36, 2003.

Po-Ling Loh and Martin J. Wainwright. Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses. *Annals of Statistics*, 41(6):3022–3049, 2013.

Cindy S. Ma, Elissa K. Deenick, Marcel Batten, and Stuart G. Tangye. The origins, function, and regulation of T follicular helper cells. *Journal of Experimental Medicine*, 209(7):1241–1253, 2012.

Georgi K Marinov, Brian A Williams, Ken McCue, Gary P Schroth, Jason Gertz, Richard M Myers, and Barbara J Wold. From single-cell to cell-pool transcriptomes: stochasticity in gene expression and RNA splicing. *Genome Research*, 24(3):496–510, mar 2014.

Florian Markowetz and Rainer Spang. Inferring cellular networks: a review. *BMC Bioinformatics*, 8(6), 2007.

Andrew McDavid, Greg Finak, Pratip K Chattopadyay, Maria Dominguez, Laurie Lam-

oreaux, Steven S Ma, Mario Roederer, and Raphael Gottardo. Data exploration, quality control and testing in single-cell qPCR-based gene expression experiments. *Bioinformatics*, 29(4):461–467, 2013.

Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34(3):1436–1462, 2006.

I. Olkin and R. F. Tate. Multivariate correlation models with mixed discrete and continuous variables. *The Annals of Mathematical Statistics*, 1961.

Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2014.

Lan V. Pham, Archito T. Tamayo, Linda C. Yoshimura, Yen Chiu Lin-Lee, and Richard J. Ford. Constitutive NF-kappaB and NFAT activation in aggressive B-cell lymphomas synergistically activates the CD154 gene and maintains lymphoma cell survival. *Blood*, 106(12):3940–3947, 2005.

Melissa L. Precopio, Michael R. Betts, Janie Parrino, David A. Price, Emma Gostick, David R. Ambrozak, Tedi E. Asher, Daniel C. Douek, Alexandre Harari, Giuseppe Pantaleo, Robert Bailer, Barney S. Graham, Mario Roederer, and Richard A. Koup. Immunization with vaccinia virus induces polyfunctional and phenotypically distinctive CD8(+) T cell responses. *Journal of Experimental Medicine*, 204(6):1405–1416, 2007.

C. Radhakirshna Rao. *Linear Statistical Inference and its Applications*. John Wiley & Sons, Ltd., 2nd edition, 1973.

Pradeep Ravikumar, Martin J. Wainwright, and John D. Lafferty. High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38 (3):1287–1319, jun 2010.

Rajen D. Shah and Richard J. Samworth. Variable selection with error control: Another look at stability selection. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 75(1):55–80, 2013.

Alex K. Shalek, Rahul Satija, Joe Shuga, John J. Trombetta, Dave Gennert, Diana Lu, Peilin Chen, Rona S. Gertner, Jellert T. Gaublomme, Nir Yosef, Schraga Schwartz, Brian Fowler, Suzanne Weaver, Jing Wang, Xiaohui Wang, Ruihua Ding, Raktima Raychowdhury, Nir Friedman, Nir Hacohen, Hongkun Park, Andrew P. May, and Aviv Regev. Single-cell RNA-seq reveals dynamic paracrine control of cellular variation. *Nature*, 510(7505):263–269, 2014.

Noah Simon and Robert Tibshirani. Standardization and the group lasso penalty. *Statistica Sinica*, 22(3):983–1001, 2012.

Wesley Tansey, Oscar Hernan Madrid Padilla, Arun Sai Suggala, and Pradeep Ravikumar. Vector-space Markov random fields via exponential families. *Proceedings of The 32nd International Conference on Machine Learning*, 37:684–692, 2015.

The Gene Ontology Consortium. Gene Ontology Consortium: going forward. *Nucleic Acids Research*, 43(D1):D1049–D1056, 2015.

Robert Tibshirani, Jacob Bien, Jerome Friedman, Trevor Hastie, Noah Simon, Jonathan Taylor, and Ryan J. Tibshirani. Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 74 (2):245–266, 2012.

Arend Voorman, Ali Shojaie, and Daniela Witten. Graph estimation with joint additive models. *Biometrika*, 101(1):85–101, 2014.

Eunho Yang, Y Baker, P Ravikumar, G Allen, and Z Liu. Mixed graphical models via exponential families. In *AISTATS 17*, volume 33, Reykjavik, Iceland, 2014.

Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 68(1): 49–67, 2006.

Ming Yuan and Yi Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94:19–35, 2007.

E-MAIL: anmcd@uw.edu

E-MAIL: rgottard@fredhutch.org

E-MAIL: nrsimon@uw.edu

E-MAIL: md5@uw.edu